Preliminary Results from a Real-World Trial of a Privacy-Preserving Crowd-Flow Sensor Network in Freiburg, Germany

Georg K.J. Fischer
Fraunhofer EMI*

Sebastian Ganter †
Fraunhofer EMI
sebastian.ganter@emi.fraunhofer.de

Uwe Benkarth
Telocate GmbH

Joan Bordoy
Telocate GmbH

Moritz Aberle
City of Freiburg

Georg Vogelbacher
Fraunhofer EMI

Tobias Bodewig
Fraunhofer EMI

Johannes Wendeberg
Telocate GmbH

Jörg Finger
Fraunhofer EMI

ABSTRACT

Estimating the crowd size and flow within inner cities and events provides crucial information for venue organizers, emergency services, and city planners. However, deploying such solutions can be a sensitive issue, since many solutions are in conflict with the GDPR. In this work, we present an architectural concept of a privacy-preserving crowd flow network based on cooperative Bluetooth Low Energy advertisements. At the core of the network is a privacy-enhancing technology, the Bloom filter, which ensures anonymization of collected data while still enabling crowd flow measurement. The sensor network has been deployed in a real-world setting, and preliminary data collected over the span of four weeks at the Christmas market in Freiburg, Germany, are presented. Besides the estimation of nearby devices, the preliminary measurement data demonstrate the feasibility of crowd flow analytics.

Keywords

Pax Counting; Passive Crowd Sensing; Privacy-Preserving Analytics; Bloom Filters; GDPR Compliance; Pedestrian Flow Estimation; Smart City Monitoring

INTRODUCTION

Events in public spaces require a high degree of planning and coordination among a multitude of stakeholders, ranging from the original event organizers and permitting bodies, such as public order offices, to emergency services. A key part of the process is determining limits on the maximum allowed occupancy of the area so that emergency evacuations can proceed in an orderly fashion. After the planning and permission phase, during the second phase in

*The Fraunhofer Institute for High-Speed Dynamics, Ernst-Mach-Institut, EMI; emi.fraunhofer.de. This work was supported by the German Federal Ministry of Research, Technology and Space (BMFTR) under grant number 13N16818 (joint project FreiburgRESIST).
†corresponding author

which the actual event takes place, these limits have to be enforced by the organizer and their contractors. The organizer is thereby required to constantly monitor the occupancy in several areas of the event.

This is mostly performed by trained security personnel who report back critical situations. Inherent to this process is the subjective judgment of whether there is a critical situation, which depends on the person themselves (their expertise and bodily parameters, such as height), as well as their point of view of the space. As limits are determined by concrete thresholds with the dimension persons/m^2, noisy and biased estimators, which such personnel inherently are, can be insufficient for this task. In Table 1, the applicable limits as defined in the respective legislation for the State of Baden-Württemberg, Germany, are given.

Table 1. Default planning densities for outdoor venues (VStättVO) to determine allowable visitor numbers

Scenario (outdoor visitor area)	Typical interpretation	Density-Limit
Seating at tables	Outdoor dining / banquet-style seating	1 person/m^2
Seating in rows	Outdoor seating in rows (e.g., open-air audience)	2 persons/m^2
Standing places (flat area)	Standing audience area (e.g., in front of a stage)	2 persons/m^2
Standing places on stepped tiers	Standing on terraces/steps (stepped standing areas)	2 persons per running meter of tier

In the planning stage, these limits have to be applied for a venue to be considered permitted by the respective authorities. However, at the onset of such venues, the actual distribution of people has not only to be monitored to enforce these limits, but also to inform emergency services in order to enable optimized responses. In this context, one of the parameters which heavily influences the time to response is the current flow of people within corridors, since these, especially in old towns, have to be traversed to reach the point of response. The envisioned system, to which this work contributes, provides for both the measurement of crowd density and the forecasting of crowd density in the near future. Both are to be displayed and analyzed in the command centers of security agencies. In particular, the intended forecasting information, which will be based on crowd flows, is designed to enable security forces to initiate measures earlier and thus 'stay ahead of the situation'. In doing so, the system is intended to provide a more comprehensive situational awareness that goes beyond selective feedback from local personnel, thereby delivering an overall picture of the situation's development.

A technical way of estimating the state of the event is therefore desirable. However, under the European General Data Protection Regulation (GDPR) and its specific national implementations, such as the Datenschutz-Grundverordnung (DSGVO) in Germany, public spaces are a high-risk context, as people cannot realistically avoid them, consent cannot usually be given, and due to the existing power imbalance between the citizen and public authorities. Following the GDPR, it is crucial that data collection is minimized and that privacy by design is applied for such systems. The legal basis for data collection is also restricted, for example when the reasons lie within another law requiring it or when it is necessary to perform a public function. Legitimate interests may also be invoked, but introduce further legal uncertainties.

Depending on the desired accuracy and costs, several technologies may be deployed in such contexts. Camera-based systems are usually the most accurate way to measure occupancy, but they also collect the highest amount of personal information, which makes justification of their deployment in public spaces particularly complex and difficult. Further, in large distributed events such as street festivals in inner cities, suitable installation points can be scarce, since position and orientation need to be tightly controlled. Similar accuracies, but with installation problems and higher costs, are offered by 2-D laser scanners, which collect point clouds of the scene. These point clouds are usually post-processed by clustering and classification algorithms to detect the number of people in a frame. Their simpler alternatives, light barriers, are wholly unsuitable, since shadowing occurs heavily in densely occupied spaces. Further, the flow of people can only be estimated within the Field-of-View (FoV), although for cameras, in principle, also between devices, which is only enabled by tracking individual people, and which in turn has strong implications for data protection.

An alternative to personal-data-sensitive devices (cameras), large-area flow estimation systems (laser scanners, thermal cameras, radar), and point-only measurements (pressure pads, light barriers, etc.) are Bluetooth/WiFi sensors (Pestalozzi et al. 2022). Since almost all persons encountered in public spaces are carrying User Equipment (UE) which support these protocols and regularly transmit discovery messages (so-called *Probe Requests* for WiFi and *Advertisements* for Bluetooth), these messages can also be used to estimate occupancy around such sensors.

However, these discovery messages contain unique device identifiers, Media Access Control (MAC) addresses, which can be regarded as personal information if a link to a concrete person can be established (see also Court of Justice of the European Union 2016). Usually, every device has several globally unique MAC addresses assigned for each network interface at the production stage. However, in recent years, device manufacturers, as well as standardization committees, have moved to employ randomized MAC addresses whenever a unique identifier is required to establish connections to unknown peers. For example, the Bluetooth Specification recommends devices to randomize their advertisement MAC address every 15 min (Bluetooth SIG 2025). This reliably breaks the link between user and device, but also leaves the possibility of flow estimation open within that time window.

In this work, we introduce the system concept of a large, deployable, and privacy-preserving crowd-flow network. Further, we detail how crowd flows can, in principle, be measured through intersections of Bloom filters, and we provide preliminary results of such a network as it was deployed at the Christmas market 2025 in Freiburg, Germany.

RELATED WORK

Bluetooth/WiFi scanners, or *Paxounters* (people counters), are widely deployed in retail (Soundararaj et al. 2020), airports (Schauer et al. 2014), and other public transport stations (Bai et al. 2017; Cheung et al. 2024; Pronello et al. 2025; Transport for London 2017), as well as in cities during festivities (Versichele et al. 2012), for general tourism analytics (Ackermann et al. 2023), at mass events (Bonne et al. 2013), and in museums (Yoshimura et al. 2016).

They are a preferred solution due to their easy deployment in temporary (events) or semi-permanent (streets, stations) settings where continuous time-series measurements are required. Use cases for situational awareness can be addressed with the measured data to identify peaks, bottlenecks, and abnormal congestion in order to support real-time decisions. In a general sense, the collected data may then be used to evaluate flow and route choice of people between zones, which simple counters (e.g., light barriers, pressure pads) cannot provide (Transport for London 2017). Due to their lightweight hardware and low installation requirements (in contrast to, e.g., vision-based systems), they can deliver cost-effective coverage at scale. Already existing WiFi access points can even be used without the need to install new hardware. Newer developments have also considered privacy-by-design implementations, since MAC addresses can be considered personal data (Rusca et al. 2024).

Privacy-Enhancing Technologies

Since crowd monitoring in public spaces cannot reasonably obtain consent, the collection of personal data is prohibited if there is no further legal basis. Strategies to handle involuntarily received personal data, such as Privacy-Enhancing Technologies (PET), are therefore necessary. When a Bluetooth or WiFi packet is sent, there is always a unique device identifier encoded within that packet, so already receiving such a packet could constitute a violation of the GDPR. Several approaches exist to enhance the privacy of the people involved:

- The baseline privacy that can be achieved with such scanners is to **minimize the data collected and introduce short retention times**. Commonly used implementations of such scanners process the MAC addresses locally and purge the collected addresses after a short time period (Klaus Wilting 2018). It should be noted that this is not a formal privacy guarantee, since, for example, in circumstances with only a few counts, it could be feasible to single out individuals.
- Collected data can directly be fed into a **pseudonymization** pipeline (Ackermann et al. 2023). The MAC addresses are hashed immediately after reception with a secret salt. This salt is rotated regularly (e.g., daily or every minute) to prevent long-term tracking (Determe et al. 2022). However, linkability can still persist, e.g., when auxiliary information is available. Hashing is generally regarded as weak anonymization (Stanciu et al. 2020).
- Apply the measure of **k-anonymity**. In this regard, unique identifiers are not only processed by hashes, which provide a 1-to-1 assignment, but are further processed in such a way that every item in the output set can be attributed to at least k items in the input set. Practically, this can be achieved by various strategies, e.g., truncating the identifiers. In turn, this lowers the detection accuracy (Stanciu et al. 2020). Bloom filters can be used in this regard to algorithmically achieve a certain level of k-anonymity while still enabling flow measurement, although with similarly decreased accuracy (Papapetrou et al. 2010; Rusca et al. 2024).

These techniques can be used when a UE is already involuntarily broadcasting personal information. This issue, however, has already been addressed by the large smartphone Operating System (OS) providers and Original Equipment Manufacturers (OEMs) by introducing MAC address randomization, starting with iOS 8 in 2014 and

with Android 6.0 in 2015 (Fenske et al. 2021). Research finds that iOS devices generally randomize their MAC addresses completely, while some Android phone manufacturers only randomize a subset of the MAC address. For WiFi MAC addresses, usually a new one is utilized for every new probe request, which makes flow measurement infeasible. For Bluetooth devices, the standard suggests a rolling randomization window of 15 min (Cäsar et al. 2022).

On the Minimum Group Size to Achieve Anonymity

For data to be regarded as anonymous, several instances of personal information have to be pooled. This, however, does not strictly speaking guarantee anonymisation in the GDPR sense. Whether data is anonymous depends on whether a person is identifiable when taking into account all means reasonably likely to be used, as well as objective factors such as costs, time, or available technology (European Parliament and the European Council 2016).

There exist several guidelines which offer practical heuristics regarding acceptable group size. The United Kingdom (UK) National Health Service (NHS), for example, speaks of suppressing small numbers (i.e., < 5) to prevent re-identification (NHS Digital 2013). The Office for National Statistics (ONS), also UK, defines the 10–5 rule, which refers to discarding counts below ten and rounding to the nearest 5 if the count is above 10 (ONS 2024). In the context of workplace statistics, the Data Protection Officer of the state of Baden-Württemberg, Germany, considers 3–5 to be too low, recommends 7, and more than 12 for sensitive contexts such as medical data (Stefan Brink 2020).

Implications for Public-Space Crowd Monitoring

The brief review shows that Bluetooth/WiFi scanners are an attractive way of measuring crowd densities and flows on a large scale in an economic manner. Use cases range from travel path optimization to enhancing situational awareness for emergency services, venue organizers, or transportation agencies. However, privacy is a major concern, since the data collected can trigger the GDPR, which is why special care has to be taken when this data is handled. Several approaches for lowering the feasibility of linkage exist, especially the deployment of k-anonymous Bloom filters, but they have not been trialed and evaluated in depth. Bluetooth in particular is a feasible candidate technology, since its randomization window is large enough to estimate flows between sensor nodes.

PRINCIPLES OF OPERATION AND SYSTEM CONCEPT

The system utilizes the Bloom filter data structure to achieve a certain degree of anonymity. In the following subsection, the fundamental working principles of a Bloom filter are detailed. Afterwards, the system concept is presented, which incorporates these principles.

A Brief Primer to Bloom Filters

The Bloom filter is a probabilistic data structure which enables fast and memory-efficient tests of whether an element x_i is part of a set or probably not. Depending on the parametrization, these properties can be tuned to achieve a certain level of k-anonymity. A Bloom filter array consists of m bits and k defined hash functions. Each hash function takes any string literal and outputs a position within the M bits. Inserting an element into the array is therefore straightforward, as shown in Figure 1. The element is hashed k times, and at each output bit position a 1 is written, leaving the rest of the array untouched. Testing is similarly straightforward: take the element to test and generate the corresponding bit position pattern. If not all positions contain a 1, the element is certainly not within the set. If the test is positive and all positions are 1, there is a chance that the element is within the set. The exact probability depends on the size of the array, the number of hash functions, and the number of already inserted elements. The more elements have been inserted into the filter, the more likely overlaps between elements become; hence, there is *plausible deniability* that the pattern could also correspond to another element. For a complete description, the reader is referred to Broder and Mitzenmacher 2004. Using, for example, the number of bits set to 1, t, one is able to calculate the most probable number of inserted elements $\hat{S}^{-1}(t)$,

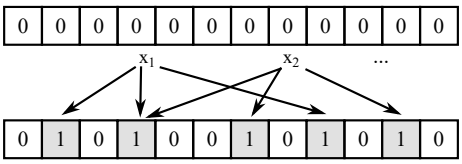

Figure 1. Illustration of inserting multiple elements x_1 and x_2 into a Bloom filter array. Due to the fact that overlap is possible when hashing the elements, there exists plausible deniability as to whether an element is part of the set or not.

$$\hat{S}^{-1}(t) = \frac{\ln\left(1 - \frac{t}{m}\right)}{k \times \ln\left(1 - \frac{1}{m}\right)}. \tag{1}$$

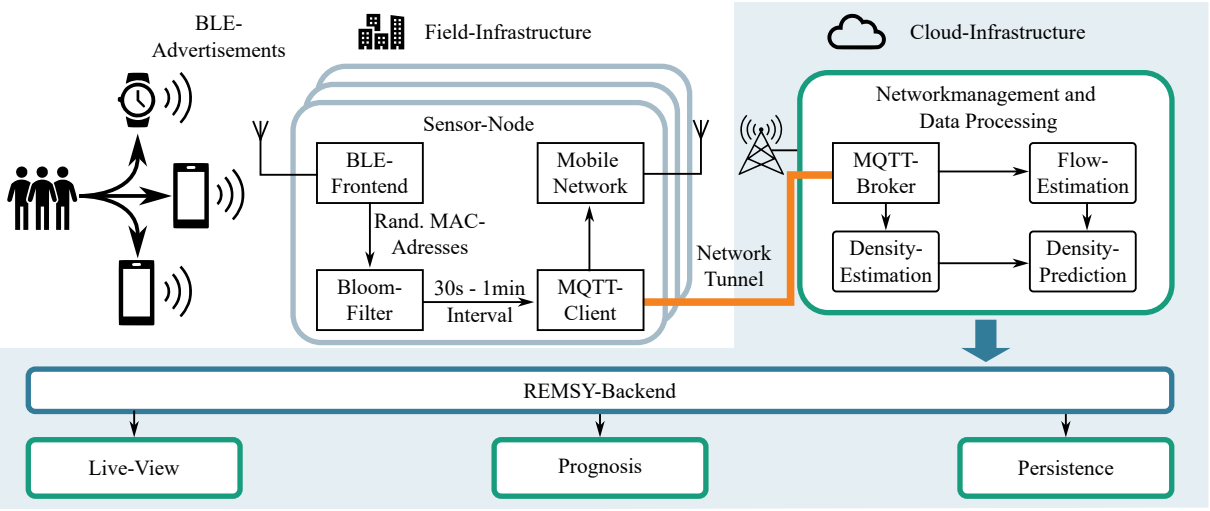

Figure 2. System concept of the Pax-counting network. Measurements are acquired in the field, and the anonymous Bloom filters are communicated to the cloud infrastructure, where data processing units such as density estimation, flow estimation, and flow prediction evaluate the data. The processed features are then made available to end users through a live view of the current state as well as forecasts, and the data is persisted.

The concrete anonymity guarantee is then called γ-K-anonymity and determines the probability γ that every inserted element has at least $K-1$ further elements which are indistinguishable from itself (Bianchi et al. 2012). To establish a certain probability, anonymization noise is usually introduced, which generates dummy elements that increase the probability of achieving K-anonymity.

System Concept

The proposed system consists of several sensor nodes deployed in the area of interest. Refer to Figure 2 for an overview. The emitted Bluetooth Low-Energy (BLE) advertisement packets of devices in proximity are filtered by their Received Signal Strength Indicator (RSSI) to establish a distance perimeter within which the density is measured. The measurement process is structured as follows:

1. An empty Bloom filter array is created, and anonymization noise is added. The Bloom filter parameters are chosen to have a 99% certainty of achieving k-anonymity greater than 30.
2. As a new advertisement packet arrives, the randomized MAC address is hashed by several hash functions and inserted into the Bloom filter.
3. This process is performed until a defined time limit of 30 s up to one minute is reached.

The Bloom filter is then sent via a cellular uplink using Message Queueing Telemetry Transport (MQTT) to an MQTT broker situated in the cloud infrastructure. There, the Bloom filter arrays of all sensor nodes are collected and processed. On the one hand, the density is estimated from the Bloom filter itself using Equation 1, and on the other hand, pedestrian flow is estimated. Local predictions are then derived by combining the flows and the measured densities.

These processed features are then fed into a Resilience Management System (REMSY), where they are consumed by the end users. The REMSY system, as part of the research project, integrates further functionality for venue organizers, emergency services, and authorities, such as planning, authorization, monitoring, and follow-up activities. The main targeted use cases in the context of crowd monitoring are monitoring density limits, route planning (for emergency services and organizers), alerts for upcoming critical densities, and the review of crowd flows in past events for lessons learned, among others.

Estimating Pedestrian Flows

To estimate pedestrian flows, different Bloom filter arrays are compared pairwise, in order to estimate how many individuals were recorded in both arrays. Thereby, the most likely number of persons included within both

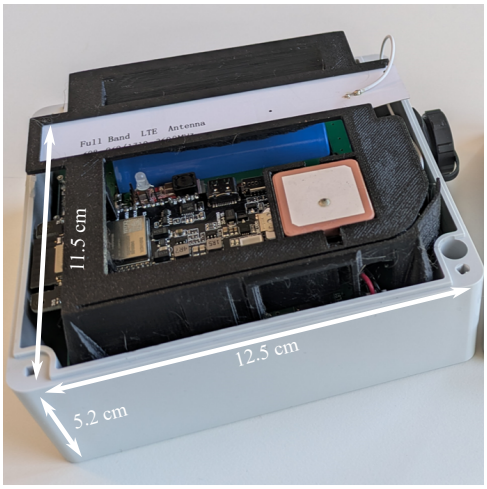

Figure 3. Photo of the deployed sensor nodes. The system integrates the BLE and cellular network RF front ends with an energy management system to bridge periods when no external power supply is available.

bloomfilter arrays is given with Eq. (2) according to Papapetrou et al. 2010.

$$\hat{S}^{-1}(t_1, t_2, t_3) = \frac{\ln\left(m - \frac{t_3 \times m - t_1 \times t_2}{m - t_1 - t_2 + t_3}\right) - \ln(m)}{k \times \ln\left(1 - \frac{1}{m}\right)} \tag{2}$$

m being the array sizes, k the number of used hash-functions and t_1, t_2 and t_3 being the Hamming weights of the individual bloomfilter arrays and their intersection, respectively. The temporal and spatial offsets between the Bloom filter recordings then allow conclusions about the direction and magnitude of pedestrian movement. By assuming a maximum measurable walking speed and specifying a minimum temporal offset, it becomes possible to determine which Bloom filter arrays should be considered for pairwise intersections. In the simplest case, the spatial offset does not represent actual walking paths but merely the geometric distance between sensor locations, so that the reachable area corresponds to a circular region, as depicted in Figure 4. The maximum allowable temporal offset between intersected Bloom filter arrays, in turn, defines the minimum non-zero person speed that can still be detected. Self-intersections of individual Bloom filter arrays are used to estimate stationary crowds that do not move at all. Simultaneously recorded Bloom filter arrays must be interpreted as redundant detections of the same individuals. The corresponding number of persons does not allow any inference about pedestrian flow, but it must be considered in the density estimation to avoid overestimating the actual person density. Finally, the resulting person flows, together with instantaneous person densities, can be used to predict short-term future densities under the assumption that variations in the overall pedestrian flow remain small.

PRELIMINARY EVALUATION

The Christmas market in Freiburg im Breisgau (Germany) runs for roughly four weeks, typically from late November through late December. In recent years it has attracted well over a million visitors overall, reaching more than 1.5 million visitors in 2025. In terms of scale, the market usually spans multiple locations in the city center and features well over 100 stalls (e.g., 127 stalls in 2025 and 140 stalls in 2024). During peak periods, especially on Advent Saturdays, daily attendance can rise to around 50,000-57,000 visitors (Franziska Pankow 2025; Laila Moscatiello 2024).

A total of 14 nodes were deployed throughout the Christmas market (see Figure 5). To evaluate the feasibility of the presented approach, we compared the estimated number of BLE devices with manual counts of people in specific areas. We conducted 13 manual counting sessions in total. During these sessions, an observer walked through a region covered by the sensors and recorded every person present using a smartphone application. The results show that the number of estimated devices correlates strongly with the manual observations ($R^2 = 0.91$), which is visible in Figure 6. However, defining the precise spatial boundaries for manual counting is challenging because nodes do not have fixed reception zones. For the area labeled *Turmstraße*, we utilized only Node 6 for validation. This node covered nearly the entire street, whereas Node 2 and Node 7 were excluded because they detected signals from devices located outside the counting area. Furthermore, to handle areas where the coverage of multiple nodes overlaps, we calculate the sum of the estimated number of elements from those nodes and then

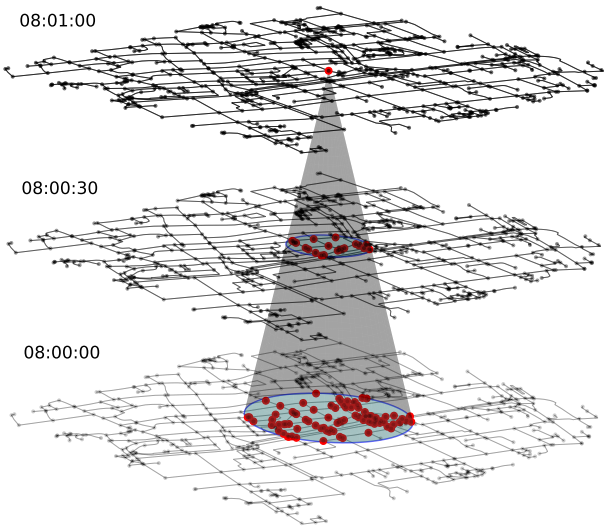

Figure 4. Estimation of reachable area for narrowing the set of bloomfilters to perform pairwise intersection.

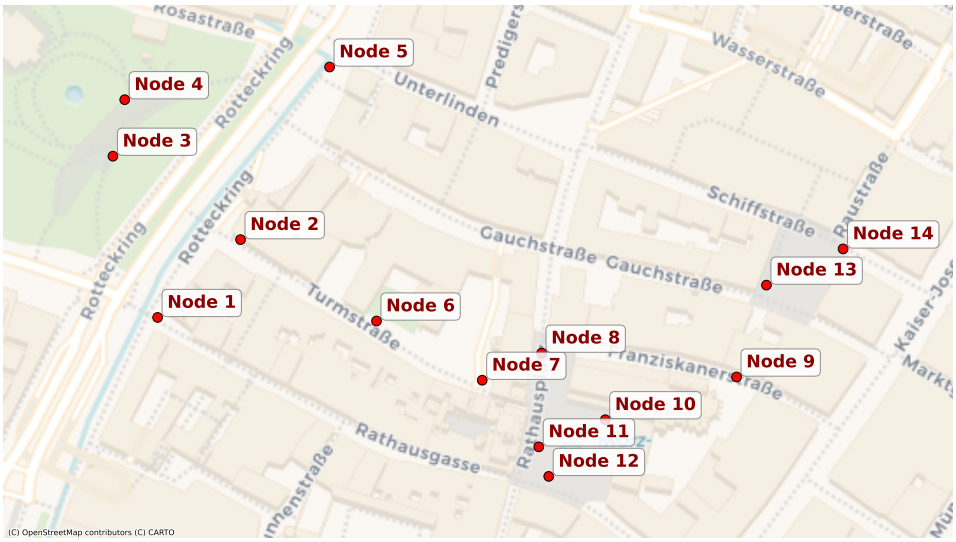

Figure 5. Location of the 14 sensor nodes deployed in the Christmas market. ***Map source: OpenStreetMaps (https://www.openstreetmap.org/copyright)***

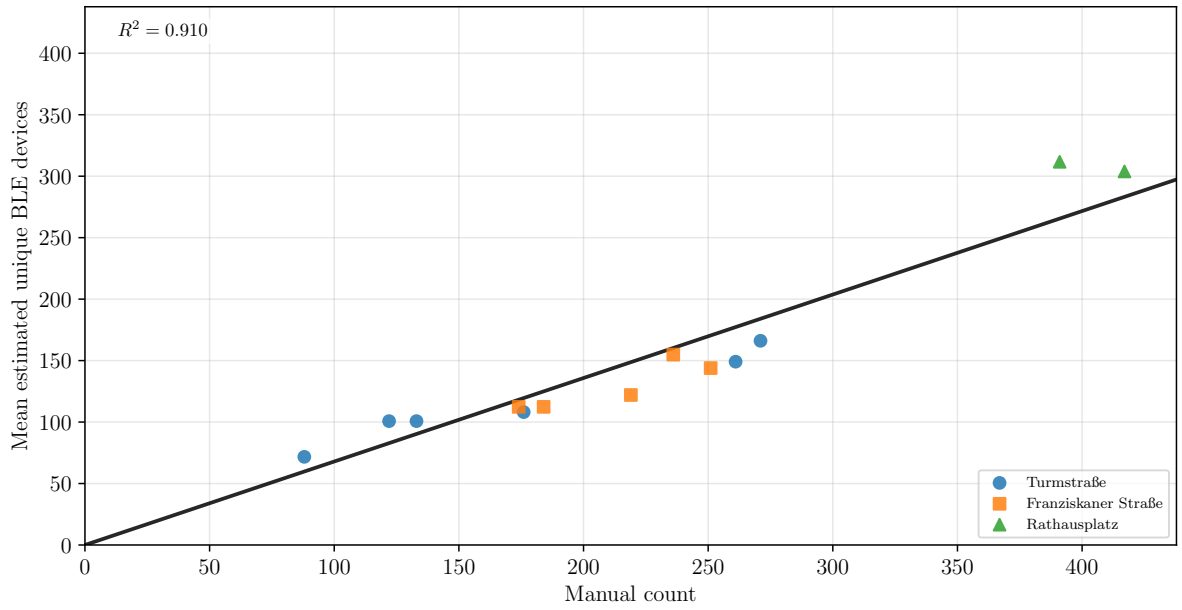

Figure 6. Correlation between manual counts and the estimated number of unique BLE devices.

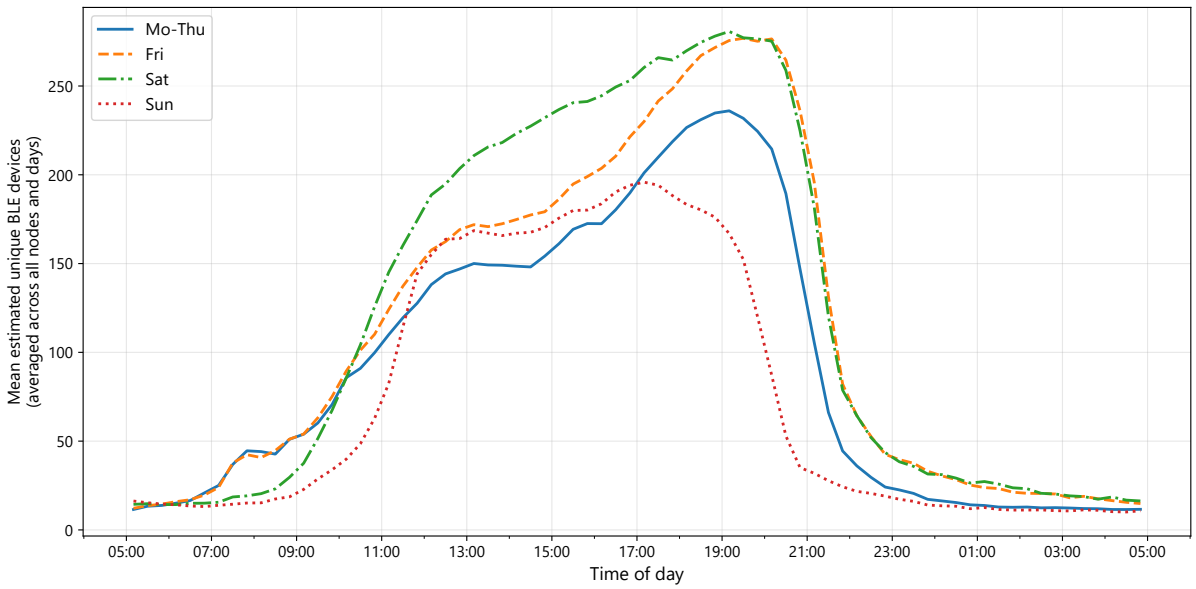

Figure 7. Average volume over all nodes. One can observe how the temporal patterns change depending on the day of the week.

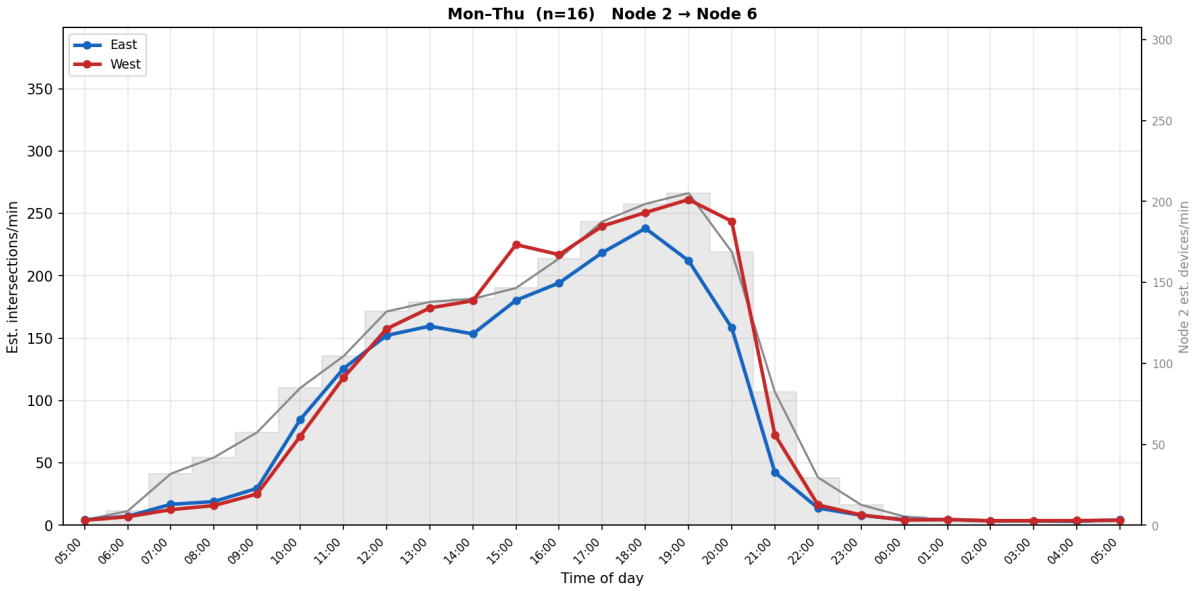

Figure 8. Bloom filter intersections between Node 2 and Node 6 from Monday to Thursday. The delay between the Bloom filters indicate the direction of the movement. A device seen in Node 2 earlier than in Node 6, indicates movement in that direction, and vice versa. We can see that the intersections increase with the volume (grey line), but there seems to be a persistent skew toward westbound movement (red line) starting at 1 p.m, peaking at 3 p.m. and again at market close at 8 p.m.

subtract the intersections. This method provides a more precise estimation of the total number of devices within a specific region, as we avoid counting the same devices multiple times.

After this initial evaluation, data were recorded continuously for 31 days, 24 hours per day, from November 21 to December 21, 2025. A first point of analysis is the variation in volume at each node across different times of day. As expected, higher concentrations of people appear around the nodes during the market's opening hours. In Figure 7, we can see that on Sundays the volume increases during the shorter opening window (11:30 to 19:30). On other days, when the market is open for longer (10:00 to 20:30), visitors tend to stay later, particularly on Fridays and Saturdays, probably due to work schedules. It is also interesting to observe a small peak around 8:00 from Monday to Friday, which is likely associated with commuting activity.

As explained earlier, Bloom filters can also be used to estimate pedestrian flows. One example is the street *Turmstraße*, which leads to the main square where most of the stands are located. Eastbound movement likely indicates people heading to the market, while westbound movement suggests people leaving it. In Fig. 8 we show the intersections between Node 2 and Node 6. If a device is observed at one node between 20 seconds and 5 minutes after being observed at the other, we infer movement in that direction. The results indicate a higher proportion of westbound movement around market close (20:30), which is consistent with people leaving the market. Since there are also stands along *Turmstraße*, it is plausible that people approach the market via other streets, and that only the closing-time pattern becomes clearly visible.

CONCLUSION AND FUTURE WORK

In this work, the use case of crowd flow monitoring is presented, as well as the regulations that make monitoring necessary. Several technologies have been discussed, showing that Bluetooth sensors can be a feasible and cost-effective solution. However, special care needs to be taken in capturing the necessary data, as it may contain personal data. For this purpose, a short literature review of possible privacy-enhancing technologies has been conducted. The notion of k-anonymity has been introduced, and the use of Bloom filter technology as a feasible PET has been detailed. The preliminary evaluation in a real-world scenario, i.e., over four weeks at a Christmas market in Germany, of the proposed sensor network not only showed the feasibility of crowd density estimation but also the possibility of crowd-flow analytics.

In upcoming research, we will systematically address the issue of balancing k-anonymity with measurement precision. From the introduced concepts, it is clear that the parameterization of the Bloom filter is the crucial part in achieving both goals. Several further experiments and longer deployments for this network are planned, in

which new parameter sets will be evaluated. The GDPR implications remain an ongoing topic, since changing the parameters always directly influences the risk of reidentification.

REFERENCES

Ackermann, L., Baum, C., Khalil, S. I., Litvin, A., and Nicklas, D. (2023). "Privacy-aware Publication of Wi-Fi Sensor Data for Crowd Monitoring and Tourism Analytics". In: *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Geo-Privacy and Data Utility for Smart Societies*. New York, NY, USA: ACM, pp. 20–23.

Bai, L., Ireson, N., Mazumdar, S., and Ciravegna, F. (2017). "Lessons learned using wi-fi and Bluetooth as means to monitor public service usage". In: *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*. Ed. by S. ". Lee, L. Takayama, and K. Truong. New York, NY, USA: ACM, pp. 432–440.

Bianchi, G., Bracciale, L., and Loreti, P. (2012). ""Better Than Nothing" Privacy with Bloom Filters: To What Extent?" In: *Privacy in statistical databases*. Ed. by J. Domingo-Ferrer. Vol. 7556. Lecture Notes in Computer Science. Berlin and Heidelberg: Springer, pp. 348–363.

Bluetooth SIG (2025). *Bluetooth Core Specification, Version 6.2*.

Bonne, B., Barzan, A., Quax, P., and Lamotte, W. (2013). "WiFiPi: Involuntary tracking of visitors at mass events". In: *2013 IEEE 14th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. IEEE, pp. 1–6.

Broder, A. and Mitzenmacher, M. (2004). "Network Applications of Bloom Filters: A Survey". In: *Internet Mathematics* 1.4, pp. 485–509.

Cäsar, M., Pawelke, T., Steffan, J., and Terhorst, G. (2022). "A survey on Bluetooth Low Energy security and privacy". In: *Computer Networks* 205, p. 108712.

Cheung, M., Cheng, Y., and Fujiyama, T. (2024). "Investigating passenger behaviour on the metro platform with Wi-Fi location tracking data: a case study of Singapore". In: *Transportation*.

Court of Justice of the European Union (2016). *Patrick Breyer v Bundesrepublik Deutschland (C-582/14)*.

Determe, J.-F., Azzagnuni, S., Singh, U., Horlin, F., and Doncker, P. de (2022). "Monitoring Large Crowds With WiFi: A Privacy-Preserving Approach". In: *IEEE Systems Journal* 16.2, pp. 2148–2159.

European Parliament and the European Council (2016). *General Data Protection Regulation (Regulation (EU) 2016/679)*.

Fenske, E., Brown, D., Martin, J., Mayberry, T., Ryan, P., and Rye, E. (2021). "Three Years Later: A Study of MAC Address Randomization In Mobile Devices And When It Succeeds". In: *Proceedings on Privacy Enhancing Technologies* 2021.3, pp. 164–181.

Franziska Pankow (2025). *Der Freiburger Weihnachtsmarkt 2025 begeistert mit neuem Besucherrekord*.

Klaus Wilting (2018). *ESP32-Paxcounter*.

Laila Moscatiello (2024). *Rund 1,33 Mio. Besucher_innen auf dem 51. Weihnachtsmarkt Freiburg*.

NHS Digital (2013). *Anonymisation Standard for Publishing Health and Social Care Data Specification*.

ONS (2024). *Comparison of post-tabular statistical disclosure control methods*.

Papapetrou, O., Siberski, W., and Nejdl, W. (2010). "Cardinality estimation and dynamic length adaptation for Bloom filters". In: *Distributed and Parallel Databases* 28.2-3, pp. 119–156.

Pestalozzi, C., Bucheli, D., and Sauter, D. (2022). *Empfehlungen zur Zählung des Fussverkehrs*. Ed. by Eidgenössisches Departement für Umwelt, Verkehr, Energie und Kommunikation UVEK.

Pronello, C., Anbarasan, D., Spoturno, F., and Terzolo, G. (2025). "A low-cost automatic people-counting system at bus stops using Wi-Fi probe requests and deep learning". In: *Public Transport* 17.1, pp. 71–100.

Rusca, R., Carluccio, A., Casetti, C., and Giaccone, P. (2024). "Privacy–preserving WiFi–based crowd monitoring". In: *Transactions on Emerging Telecommunications Technologies* 35.3.

Schauer, L., Werner, M., and Marcus, P. (2014). "Estimating Crowd Densities and Pedestrian Flows Using Wi-Fi and Bluetooth". In: *Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*. Ed. by M. Youssef. ICST.

Soundararaj, B., Cheshire, J., and Longley, P. (2020). "Estimating real-time high-street footfall from Wi-Fi probe requests". In: *International Journal of Geographical Information Science* 34.2, pp. 325–343.

Stanciu, V.-D., van Steen, M., Dobre, C., and Peter, A. (2020). "k-Anonymous Crowd Flow Analytics". In: *MobiQuitous 2020 - 17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*. Ed. by M. Mühlhäuser, G. C. Polyzos, F. Michahelles, A. S. Guinea, and L. Wang. New York, NY, USA: ACM, pp. 376–385.

Stefan Brink (2020). *Arbeitnehmerdatenschutz: Zwischen wirtschaftlicher Abhängigkeit und informationeller Selbstbestimmung*.

Transport for London (2017). *Review of the TfL WiFi pilot*.

Versichele, M., Neutens, T., Delafontaine, M., and van de Weghe, N. (2012). "The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities". In: *Applied Geography* 32.2, pp. 208–220.

Yoshimura, Y., Krebs, A., and Ratti, C. (2016). "An analysis of visitors' length of stay through noninvasive Bluetooth monitoring in the Louvre Museum". In.